arXiv:1411.0073v1 [stat.ML] 1 Nov 2014

Learning Mixed Multinomial Logit Model from Ordinal Data

Sewoong Oh
Dept. of Industrial and Enterprise Systems Engr.
University of Illinois at Urbana-Champaign
Urbana, IL 61801
swoh@illinois.edu

Devavrat Shah
Department of Electrical Engineering
Massachussetts Institute of Technology
Cambridge, MA 02139
devavrat@mit.edu

Abstract

Motivated by generating personalized recommendations using ordinal (or preference) data, we study the question of learning a mixture of MultiNomial Logit (MNL) model, a parameterized class of distributions over permutations, from partial ordinal or preference data (e.g. pair-wise comparisons). Despite its long standing importance across disciplines including social choice, operations research and revenue management, little is known about this question. In case of single MNL models (no mixture), computationally and statistically tractable learning from pair-wise comparisons is feasible. However, even learning mixture with two MNL components is infeasible in general.

Given this state of affairs, we seek conditions under which it is feasible to learn the mixture model in both computationally and statistically efficient manner. We present a sufficient condition as well as an efficient algorithm for learning mixed MNL models from partial preferences/comparisons data. In particular, a mixture of r MNL components over n objects can be learnt using samples whose size scales polynomially in n and r (concretely, $r^{3.5}n^3(\log n)^4$, with $r \ll n^{2/7}$ when the model parameters are sufficiently *incoherent*). The algorithm has two phases: first, learn the pair-wise marginals for each component using tensor decomposition; second, learn the model parameters for each component using RANKCENTRALITY introduced by Negahban et al. In the process of proving these results, we obtain a generalization of existing analysis for tensor decomposition to a more realistic regime where only partial information about each sample is available.

1 Introduction

Background. Popular recommendation systems such as collaborative filtering are based on a partially observed *ratings* matrix. The underlying hypothesis is that the true/latent score matrix is low-rank and we observe its partial, noisy version. Therefore, matrix completion algorithms are used for learning, cf. [8, 14, 15, 20]. In reality, however, observed preference data is not just scores. For example, clicking one of the many choices while browsing provides partial order between clicked choice versus other choices. Further, scores do convey ordinal information as well, e.g. score of 4 for paper A and score of 7 for paper B by a reviewer suggests ordering B $>$ A. Similar motivations led Samuelson to propose the *Axiom of revealed preference* [21] as the model for rational behavior. In a nutshell, it states that *consumers have latent order of all objects, and the revealed preferences through actions/choices are consistent with this order.* If indeed all consumers had identical ordering, then learning preference from partial preferences is effectively the question of sorting.

In practice, individuals have different orderings of interest, and further, each individual is likely to make noisy choices. This naturally suggests the following model – each individual has a latent distribution over orderings of objects of interest, and the revealed partial preferences are consistent

with it, i.e. samples from the distribution. Subsequently, the preference of the population as a whole can be associated with a distribution over permutations. Recall that the low-rank structure for score matrices, as a model, tries to capture the fact that there are only a few different types of choice profile. In the context of modeling consumer choices as distribution over permutation, MultiNomial Logit (MNL) model with a small number of mixture components provides such a model.

Mixed MNL. Given n objects or choices of interest, an MNL model is described as a parametric distribution over permutations of n with parameters $\mathbf{w} = [w_i] \in \mathbb{R}^n$: each object $i,\ 1 \leq i \leq n$, has a parameter $w_i > 0$ associated with it. Then the permutations are generated randomly as follows: choose one of the n objects to be ranked 1 at random, where object i is chosen to be ranked 1 with probability $w_i/(\sum_{j=1}^{n} w_j)$. Let i_1 be object chosen for the first position. Now to select second ranked object, choose from remaining with probability proportional to their weight. We repeat until all objects for all ranked positions are chosen. It can be easily seen that, as per this model, an item i is ranked higher than j with probability $w_i/(w_i + w_j)$.

In the mixed MNL model with $r \geq 2$ mixture components, each component corresponds to a different MNL model: let $\mathbf{w}^{(1)}, \ldots, \mathbf{w}^{(r)}$ be the corresponding parameters of the r components. Let $\mathbf{q} = [q_a] \in [0,1]^r$ denote the mixture distribution, i.e. $\sum_a q_a = 1$. To generate a permutation at random, first choose a component $a \in \{1, \ldots, r\}$ with probability q_a, and then draw random permutation as per MNL with parameters $\mathbf{w}^{(a)}$.

Brief history. The MNL model is an instance of a class of models introduced by Thurstone [23]. The description of the MNL provided here was formally established by McFadden [17]. The same model (in form of pair-wise marginals) was introduced by Zermelo [25] as well as Bradley and Terry [7] independently. In [16], Luce established that MNL is the only distribution over permutation that satisfies the axiom of Independence from Irrelevant Alternatives.

On learning distributions over permutations, the question of learning single MNL model and more generally instances of Thurstone's model have been of interest for quite a while now. The maximum likelihood estimator, which is logistic regression for MNL, has been known to be consistent in large sample limit, cf. [13]. Recently, RANKCENTRALITY [19] was established to be statistical efficient. For learning sparse mixture model, i.e. distribution over permutations with each mixture being delta distribution, [11] provided sufficient conditions under which mixtures can be learnt exactly using pair-wise marginals – effectively, as long as the number of components scaled as $o(\log n)$ where components satisfied appropriate *incoherence* condition, a simple iterative algorithm could recover the mixture. However, it is not robust with respect to noise in data or finite sample error in marginal estimation. Other approaches have been proposed to recover model using convex optimization based techniques, cf. [10, 18]. MNL model is a special case of a larger family of discrete choice models known as the Random Utility Model (RUM), and an efficient algorithm to learn RUM is introduced in [22]. Efficient algorithms for learning RUMs from partial rankings has been introduced in [3, 4]. We note that the above list of references is very limited, including only closely related literature. Given the nature of the topic, there are a lot of exciting lines of research done over the past century and we shall not be able to provide comprehensive coverage due to a space limitation.

Problem. Given observations from the mixed MNL, we wish to learn the model parameters, the mixing distribution $\mathbf{q}$, and parameters of each component $\mathbf{w}^{(1)}, \ldots, \mathbf{w}^{(r)}$. The observations are in form of pair-wise comparisons. Formally, to generate an observation, first one of the r mixture components is chosen; and then for ℓ of all possible $\binom{n}{2}$ pairs, comparison outcome is observed as per this MNL component[1]. These ℓ pairs are chosen, uniformly at random, from a pre-determined $N \leq \binom{n}{2}$ pairs: $\{(i_k, j_k), 1 \leq k \leq N\}$. We shall assume that the selection of N is such that the undirected graph $G = ([n], E)$, where $E = \{(i_k, j_k) : 1 \leq k \leq N\}$, is connected.

We ask following questions of interest: Is it always feasible to learn mixed MNL? If not, under what conditions and how many samples are needed? How computationally expensive are the algorithms?

[1]We shall assume that, outcomes of these ℓ pairs are independent of each other, but coming from the same MNL mixture component. This is effectively true even they were generated by first sampling a permutation from the chosen MNL mixture component, and then observing implication of this permutation for the specific ℓ pairs, as long as they are distinct due to the Irrelevance of Independent Alternative hypothesis of Luce that is satisfied by MNL.

We briefly recall a recent result [1] that suggests that it is impossible to learn mixed MNL models in general. One such example is described in Figure 1. It depicts an example with $n = 4$ and $r = 2$ and a uniform mixture distribution. For the first case, in mixture component 1, with probability 1 the ordering is $a > b > c > d$ (we denote $n = 4$ objects by a, b, c and d); and in mixture component 2, with probability 1 the ordering is $b > a > d > c$. Similarly for the second case, the two mixtures are made up of permutations $b > a > c > d$ and $a > b > d > c$. It is easy to see the distribution over any 3-wise comparisons generated from these two mixture models is identical. Therefore, it is impossible to differentiate these two using 3-wise or pair-wise comparisons. In general, [1] established that there exist mixture distributions with $r \leq n/2$ over n objects that are impossible to distinguish using $\log n$-wise comparison data. That is, learning mixed MNL is not always possible.

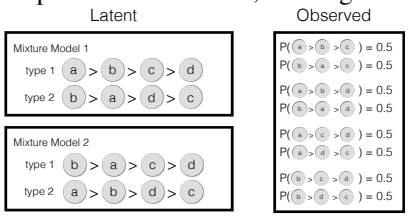
Latent
Observed
Mixture Model 1
type 1 a > b > c > d
type 2 b > a > d > c
Mixture Model 2
type 1 b > a > c > d
type 2 a > b > d > c
P(a > b > c) = 0.5
P(b > a > c) = 0.5
P(a > b > d) = 0.5
P(b > a > d) = 0.5
P(a > c > d) = 0.5
P(a > d > c) = 0.5
P(b > c > d) = 0.5
P(b > d > c) = 0.5

Figure 1: Two mixture models that cannot be differentiated even with 3-wise preference data.

Contributions. The main contribution of this work is identification of sufficient conditions under which mixed MNL model can be learnt efficiently, both statistically and computationally. Concretely, we propose a two-phase learning algorithm: in the first phase, using a tensor decomposition method for learning mixture of discrete product distribution, we identify pair-wise marginals associated with each of the mixture; in the second phase, we use these pair-wise marginals associated with each mixture to learn the parameters associated with each of the MNL mixture component.

The algorithm in the first phase builds upon the recent work by Jain and Oh [12]. In particular, Theorem 3 generalizes their work for the setting where for each sample, we have limited information - as per [12], we would require that each individual gives the entire permutation; instead, we have extended the result to be able to cope with the current setting when we only have information about ℓ, potentially finite, pair-wise comparisons. The algorithm in the second phase utilizes RANK-CENTRALITY [19]. Its analysis in Theorem 4 works for setting where observations are no longer independent, as required in [19].

We find that as long as certain rank and *incoherence* conditions are satisfied by the parameters of each of the mixture, the above described two phase algorithm is able to learn mixture distribution $\mathbf{q}$ and parameters associated with each mixture, $\mathbf{w}^{(1)}, \ldots, \mathbf{w}^{(r)}$ faithfully using samples that scale polynomially in n and r – concretely, the number of samples required scale as $r^{3.5}n^3(\log n)^4$ with constants dependent on the incoherence between mixture components, and as long as $r \ll n^{2/7}$ as well as G, the graph of potential comparisons, is a spectral expander with the total number of edges scaling as $N = O(n \log n)$. For the precise statement, we refer to Theorem 1.

The algorithms proposed are iterative, and primarily based on spectral properties of underlying tensors/matrices with provable, fast convergence guarantees. That is, algorithms are not only polynomial time, they are practical enough to be scalable for high dimensional data sets.

Notations. We use $[N] = \{1, \ldots, N\}$ for the first N positive integers. We use $\otimes$ to denote the outer product such that $(x \otimes y \otimes z)_{ijk} = x_i y_j z_k$. Given a third order tensor $T \in \mathbb{R}^{n_1 \times n_2 \times n_3}$ and a matrix $U \in \mathbb{R}^{n_1 \times r_1}, V \in \mathbb{R}^{n_2 \times r_2}, W \in \mathbb{R}^{n_3 \times r_3}$, we define a linear mapping $T[U, V, W] \in \mathbb{R}^{r_1 \times r_2 \times r_3}$ as $T[U, V, W]_{abc} = \sum_{i,j,k} T_{ijk} U_{ia} V_{jb} W_{kc}$. We let $\|x\| = \sqrt{\sum_i x_i^2}$ be the Euclidean norm of a vector, $\|M\|_2 = \max_{\|x\| \leq 1, \|y\| \leq 1} x^T M y$ be the operator norm of a matrix, and $\|M\|_F = \sqrt{\sum_{i,j} M_{ij}^2}$ be the Frobenius norm. We say an event happens with high probability (w.h.p) if the probability is lower bounded by $1 - f(n)$ such that $f(n) = o(1)$ as n scales to ∞.

2 Main result

In this section, we describe the main result: sufficient conditions under which mixed MNL models can be learnt using tractable algorithms. We provide a useful illustration of the result as well as discuss its implications.

Definitions. Let $\mathcal{S}$ denote the collection of observations, each of which is denoted as N dimensional, $\{-1, 0, +1\}$ valued vector. Recall that each observation is obtained by first selecting one of the r mixture MNL component, and then viewing outcomes, as per the chosen MNL mixture component, of ℓ randomly chosen pair-wise comparisons from the N pre-determined comparisons $\{(i_k, j_k) : 1 \leq i_k \neq j_k \leq n, 1 \leq k \leq N\}$. Let $x_t \in \{-1, 0, +1\}^N$ denote the tth observation with $x_{t,k} = 0$ if the kth pair (i_k, j_k) is not chosen amongst the ℓ randomly chosen pairs, and $x_{t,k} = +1$ (respectively -1) if $i_k < j_k$ (respectively $i_k > j_k$) as per the chosen MNL mixture component. By definition, it is easy to see that for any $t \in \mathcal{S}$ and $1 \leq k \leq N$,

$$\mathbb{E}[x_{t,k}] = \frac{\ell}{N}\Big[\sum_{a=1}^{r} q_a P_{ka}\Big], \text{ where } P_{ka} = \frac{w_{j_k}^{(a)} - w_{i_k}^{(a)}}{w_{j_k}^{(a)} + w_{i_k}^{(a)}}. \tag{1}$$

We shall denote $P_a = [P_{ka}] \in [-1, 1]^N$ for $1 \leq a \leq r$. Therefore, in a vector form

$$\mathbb{E}[x_t] = \frac{\ell}{N} P\mathbf{q}, \text{ where } P = [P_1 \dots P_r] \in [-1, 1]^{N \times r} . \tag{2}$$

That is, P is a matrix with r columns, each representing one of the r mixture components and $\mathbf{q}$ is the mixture probability. By independence, for any $t \in \mathcal{S}$, and any two different pairs $1 \leq k \neq m \leq N$,

$$\mathbb{E}[x_{t,k} x_{t,m}] = \frac{\ell^2}{N^2}\Big[\sum_{a=1}^{r} q_a P_{ka} P_{ma}\Big]. \tag{3}$$

Therefore, the $N \times N$ matrix $\mathbb{E}[x_t x_t^T]$ or equivalently tensor $\mathbb{E}[x_t \otimes x_t]$ is proportional to M_2 except in diagonal entries, where

$$M_2 = PQP^T \equiv \sum_{a=1}^{r} q_a (P_a \otimes P_a) , \tag{4}$$

$Q = \mathrm{diag}(\mathbf{q})$ being diagonal matrix with its entries being mixture probabilities, $\mathbf{q}$. In a similar manner, the tensor $\mathbb{E}[x_t \otimes x_t \otimes x_t]$ is proportional to M_3 (except in $O(N^2)$ entries), where

$$M_3 = \sum_{a=1}^{r} q_a (P_a \otimes P_a \otimes P_a). \tag{5}$$

Indeed, empirical estimates $\hat{M}_2$ and $\hat{M}_3$, defined as

$$\hat{M}_2 = \frac{1}{|\mathcal{S}|}\Big[\sum_{t \in \mathcal{S}} x_t \otimes x_t\Big], \text{ and } \hat{M}_3 = \frac{1}{|\mathcal{S}|}\Big[\sum_{t \in \mathcal{S}} x_t \otimes x_t \otimes x_t\Big], \tag{6}$$

provide good proxy for M_2 and M_3 for large enough number of samples; and shall be utilized crucially for learning model parameters from observations.

Sufficient conditions for learning. With the above discussion, we state sufficient conditions for learning the mixed MNL in terms of properties of M_2:

C1. M_2 has rank r; let $\sigma_1(M_2), \sigma_r(M_2) > 0$ be the largest and smallest singular values of M_2.

C2. For a large enough universal constant $C' > 0$,

$$N \geq C' r^{3.5}\, \mu^6(M_2) \Big(\frac{\sigma_1(M_2)}{\sigma_r(M_2)}\Big)^{4.5}. \tag{7}$$

In the above, $\mu(M_2)$ represents incoherence of a symmetric matrix M_2. We recall that for a symmetric matrix $M \in \mathbb{R}^{N \times N}$ of rank r with singular value decomposition $M = USU^T$, the incoherence is defined as

$$\mu(M) = \sqrt{\frac{N}{r}}\Big(\max_{i \in [N]} \|U_i\|\Big). \tag{8}$$

C3. The undirected graph $G = ([n], E)$ with $E = \{(i_k, j_k) : 1 \leq k \leq N\}$ is connected. Let $A \in \{0, 1\}^{n \times n}$ be adjacency matrix with $A_{ij} = 1$ if $(i, j) \in E$ and 0 otherwise; let $D = \mathrm{diag}(d_i)$ with d_i being degree of vertex $i \in [n]$ and let $L_G = D^{-1}A$ be normalized Laplacian of G. Let $d_{\max} = \max_i d_i$ and $d_{\min} = \min_i d_i$. Let the n eigenvalues of stochastic matrix L_G be $1 = \lambda_1(L_G) \geq \dots \lambda_n(L_G) \geq -1$. Define spectral gap of G:

$$\xi(G) = 1 - \max\{\lambda_2(L), -\lambda_n(L)\}. \tag{9}$$

Note that we choose a graph $G = ([n], E)$ to collect pairwise data on, and we want to use a graph that is connected, has a large spectral gap, and has a small number of edges. In condition (C3), we need connectivity since we cannot estimate the relative strength between disconnected components (e.g. see [13]). Further, it is easy to generate a graph with spectral gap $\xi(G)$ bounded below by a universal constant (e.g. $1/100$) and the number of edges $N = O(n \log n)$, for example using the configuration model for Erdös-Renyi graphs. In condition (C2), we require the matrix M_2 to be sufficiently incoherent with bounded $\sigma_1(M_2)/\sigma_r(M_2)$. For example, if $q_{\max}/q_{\min} = O(1)$ and the profile of each type in the mixture distribution is sufficiently different, i.e. $\langle P_a, P_b \rangle / (\|P_a\| \|P_b\|) < 1/(2r)$, then we have $\mu(M_2) = O(1)$ and $\sigma_1(M_2)/\sigma_r(M_2) = O(1)$. We define $b = \max_{a=1}^{r} \max_{i,j \in [n]} w_i^{(a)}/w_j^{(a)}$, $q_{\max} = \max_a q_a$, and $q_{\min} = \min_a q_a$. The following theorem provides a bound on the error and we refer to the appendix for a proof.

Theorem 1. *Consider a mixed MNL model satisfying conditions (C1)-(C3). Then for any $\delta \in (0,1)$, there exists positive numerical constants C, C' such that for any positive ε satisfying*

$$0 < \varepsilon < \Big(\frac{q_{\min} \xi^2(G) d_{\min}^2}{16 q_{\max}\, r\, \sigma_1(M_2) b^5 d_{\max}^2} \Big)^{0.5}, \tag{10}$$

Algorithm 1 produces estimates $\hat{\mathbf{q}} = [\hat{q}_a]$ and $\hat{\mathbf{w}} = [\hat{\mathbf{w}}^{(a)}]$ so that with probability at least $1 - \delta$,

$$\begin{aligned} \big|\, \hat{q}_a - q_a \,\big| &\le \varepsilon, \text{ and} \\ \frac{\|\hat{\mathbf{w}}^{(a)} - \mathbf{w}^{(a)}\|}{\|\mathbf{w}^{(a)}\|} &\le C \Big(\frac{r\, q_{\max}\, \sigma_1(M_2) b^5 d_{\max}^2}{q_{\min} \xi^2(G) d_{\min}^2} \Big)^{0.5} \varepsilon, \end{aligned} \tag{11}$$

for all $a \in [r]$, as long as

$$|\mathcal{S}| \ge C' \frac{r N^4 \log(N/\delta)}{q_{\min} \sigma_1(M_2)^2 \varepsilon^2} \Big(\frac{1}{\ell^2} + \frac{\sigma_1(M_2)}{\ell N} + \frac{r^4 \sigma_1(M_2)^4}{\sigma_r(M_2)^5} \Big). \tag{12}$$

An illustration of Theorem 1. To understand the applicability of Theorem 1, consider a concrete example with $r = 2$; let the corresponding weights $\mathbf{w}^{(1)}$ and $\mathbf{w}^{(2)}$ be generated by choosing each weight uniformly from $[1, 2]$. In particular, the rank order for each component is a uniformly random permutation. Let the mixture distribution be uniform as well, i.e. $\mathbf{q} = [0.5\ 0.5]$. Finally, let the graph $G = ([n], E)$ be chosen as per the Erdös-Rényi model with each edge chosen to be part of the graph with probability $\bar{d}/n$, where $\bar{d} > \log n$. For this example, it can be checked that Theorem 1 guarantees that for $\varepsilon \le C/\sqrt{n\bar{d}}$, $|\mathcal{S}| \ge C' n^2 \bar{d}^2 \log(n\bar{d}/\delta)/(\ell \varepsilon^2)$, and $n\bar{d} \ge C'$, we have for all $a \in \{1, 2\}$, $|\hat{q}_a - q_a| \le \varepsilon$ and $\|\hat{\mathbf{w}}^{(a)} - \mathbf{w}^{(a)}\|/\|\mathbf{w}^{(a)}\| \le C'' \sqrt{n\bar{d}}\, \varepsilon$. That is, for $\ell = \Theta(1)$ and choosing $\varepsilon = \varepsilon'/(\sqrt{n\bar{d}})$, we need sample size of $|\mathcal{S}| = O(n^3 \bar{d}^3 \log n)$ to guarantee error in both $\hat{\mathbf{q}}$ and $\hat{\mathbf{w}}$ smaller than ε'. Instead, if we choose $\ell = \Theta(n\bar{d})$, we only need $|\mathcal{S}| = O((n\bar{d})^2 \log n)$. Limited samples per observation leads to penalty of factor of $(n\bar{d}/\ell)$ in sample complexity. To provide bounds on the problem parameters for this example, we use standard concentration arguments. It is well known for Erdös-Rényi random graphs (see [6]) that, with high probability, the number of edges concentrates in $[(1/2)\bar{d}\, n, (3/2)\bar{d}\, n]$ implying $N = \Theta(\bar{d}n)$, and the degrees also concentrate in $[(1/2)\bar{d}, (3/2)\bar{d}]$, implying $d_{\max} = d_{\min} = \Theta(\bar{d})$. Also using standard concentration arguments for spectrum of random matrices, it follows that the spectral gap of G is bounded by $\xi \ge 1 - (C/\sqrt{\bar{d}}) = \Theta(1)$ w.h.p. Since we assume the weights to be in $[1, 2]$, the dynamic range is bounded by $b \le 2$. The following Proposition shows that $\sigma_1(M_2) = \Theta(N) = \Theta(\bar{d}n)$, $\sigma_2(M_2) = \Theta(\bar{d}n)$, and $\mu(M_2) = \Theta(1)$.

Proposition 2.1. *For the above example, when $\bar{d} \ge \log n$, $\sigma_1(M_2) \le 0.02N$, $\sigma_2(M_2) \ge 0.017N$, and $\mu(M_2) \le 15$ with high probability.*

Supposen now for general r, we are interested in well-behaved scenario where $q_{\max} = \Theta(1/r)$ and $q_{\min} = \Theta(1/r)$. To achieve arbitrary small error rate for $\|\hat{\mathbf{w}}^{(a)} - \mathbf{w}^{(a)}\|/\|\mathbf{w}^{(a)}\|$, we need $\epsilon = O(1/\sqrt{r\, N})$, which is achieved by sample size $|\mathcal{S}| = O(r^{3.5} n^3 (\log n)^4)$ with $\bar{d} = \log n$.

3 Algorithm

We describe the algorithm achieving the bound in Theorem 1. Our approach is two-phased. First, learn the moments for mixtures using a tensor decomposition, cf. Algorithm 2: for each type $a \in [r]$,

produce estimate $\hat{q}_a \in \mathbb{R}$ of the mixture weight q_a and estimate $\hat{P}_a = [\hat{P}_{1a} \ \dots \ \hat{P}_{Na}]^T \in \mathbb{R}^N$ of the expected outcome $P_a = [P_{1a} \ \dots \ P_{Na}]^T$ defined as in (1). Secondly, for each a, using the estimate $\hat{P}_a$, apply RANKCENTRALITY, cf. Section 3.2, to estimate $\hat{\mathbf{w}}^{(a)}$ for the MNL weights $\mathbf{w}^{(a)}$.

Algorithm 1

Input: Samples $\{x_t\}_{t\in\mathcal{S}}$, number of types r, number of iterations T_1, T_2, graph $G([n], E)$
$\{(\hat{q}_a, \hat{P}_a)\}_{a\in[r]} \leftarrow$ SPECTRALDIST $(\{x_t\}_{t\in\mathcal{S}}, r, T_1)$ (see Algorithm 2)
for $a = 1, \dots, r$ **do**
  set $\tilde{P}_a \leftarrow \mathcal{P}_{[-1,1]}(\hat{P}_a)$ where $\mathcal{P}_{[-1,1]}(\cdot)$ is the projection onto $[-1, 1]^N$
  $\hat{w}^{(a)} \leftarrow$ RANKCENTRALITY $\Big(G, \tilde{P}_a, T_2\Big)$ (see Section 3.2)
end for
Output: $\{(\hat{q}^{(a)}, \hat{\mathbf{w}}^{(a)})\}_{a\in[r]}$

To achieve Theorem 1, $T_1 = \Theta\big(\log(N\,|\mathcal{S}|)\big)$ and $T_2 = \Theta\big(\, b^2 d_{\max}(\log n + \log(1/\varepsilon))/(\xi d_{\min})\,\big)$ is sufficient. Next, we describe the two phases of algorithms and associated technical results.

3.1 Phase 1: Spectral decomposition.

To estimate P and $\mathbf{q}$ from the samples, we shall use tensor decomposition of $\hat{M}_2$ and $\hat{M}_3$, the empirical estimation of M_2 and M_3 respectively, recall (4)-(6). Let $M_2 = U_{M_2}\Sigma_{M_2}U_{M_2}^T$ be the eigenvalue decomposition and let

$$H \ = \ M_3[U_{M_2}\Sigma_{M_2}^{-1/2}, U_{M_2}\Sigma_{M_2}^{-1/2}, U_{M_2}\Sigma_{M_2}^{-1/2}] \ .$$

The next theorem shows that M_2 and M_3 are sufficient to learn P and $\mathbf{q}$ exactly, when M_2 has rank r (throughout, we assume that $r \ll n \leq N$).

Theorem 2 (Theorem 3.1 [12])**.** *Let $M_2 \in \mathbb{R}^{N\times N}$ have rank r. Then there exists an orthogonal matrix $V^H = [v_1^H \ v_2^H \ \dots \ v_r^H] \in \mathbb{R}^{r\times r}$ and eigenvalues λ_a^H, $1 \leq a \leq r$, such that the orthogonal tensor decomposition of H is*

$$H \ = \ \sum_{a=1}^{r} \lambda_a^H (v_a^H \otimes v_a^H \otimes v_a^H).$$

Let $\Lambda^H = \mathrm{diag}(\lambda_1^H, \dots, \lambda_r^H)$. Then the parameters of the mixture distribution are

$$P \ = \ U_{M_2}\Sigma_{M_2}^{1/2}V^H\Lambda^H \ \ \textit{and} \quad Q \ = \ (\Lambda^H)^{-2} \ .$$

The main challenge in estimating M_2 (resp. M_3) from empirical data are the diagonal entires. In [12], alternating minimization approach is used for matrix completion to find the missing diagonal entries of M_2, and used a least squares method for estimating the tensor H directly from the samples. Let Ω_2 denote the set of off-diagonal indices for an $N \times N$ matrix and Ω_3 denote the off-diagonal entries of an $N \times N \times N$ tensor such that the corresponding projections are defined as

$$\mathcal{P}_{\Omega_2}(M)_{ij} \ \equiv \ \begin{cases} M_{ij} & \text{if } i \neq j \ , \\ 0 & \text{otherwise} \ . \end{cases} \quad \text{and} \quad \mathcal{P}_{\Omega_3}(T)_{ijk} \ \equiv \ \begin{cases} T_{ijk} & \text{if } i \neq j, j \neq k, k \neq i \ , \\ 0 & \text{otherwise} \ . \end{cases}$$

for $M \in \mathbb{R}^{N\times N}$ and $T \in \mathbb{R}^{N\times N\times N}$.

In lieu of above discussion, we shall use $\mathcal{P}_{\Omega_2}\big(\hat{M}_2\big)$ and $\mathcal{P}_{\Omega_3}\big(\hat{M}_3\big)$ to obtain estimation of diagonal entries of M_2 and M_3 respectively. To keep technical arguments simple, we shall use first $|\mathcal{S}|/2$ samples based $\hat{M}_2$, denoted as $\hat{M}_2\big(1, \frac{|\mathcal{S}|}{2}\big)$ and second $|\mathcal{S}|/2$ samples based $\hat{M}_3$, denoted by $\hat{M}_3\big(\frac{|\mathcal{S}|}{2}+1, |\mathcal{S}|\big)$ in Algorithm 2.

Next, we state correctness of Algorithm 2 when $\mu(M_2)$ is small; proof is in Appendix.

Theorem 3. *There exists universal, strictly positive constants $C, C' > 0$ such that for all $\varepsilon \in (0, C)$ and $\delta \in (0, 1)$, if*

$$\begin{aligned} |\mathcal{S}| \ &\geq \ C' \frac{rN^4 \log(N/\delta)}{q_{\min}\sigma_1(M_2)^2\varepsilon^2}\Big(\frac{1}{\ell^2} + \frac{\sigma_1(M_2)}{\ell N} + \frac{r^4\sigma_1(M_2)^4}{\sigma_r(M_2)^5}\Big) \ , \ \textit{and} \\ N \ &\geq \ C' r^{3.5}\mu^6 \Big(\frac{\sigma_1(M_2)}{\sigma_r(M_2)}\Big)^{4.5} \ , \end{aligned}$$

Algorithm 2 SPECTRALDIST**:** Moment method for Mixture of Discrete Distribution [12]

1: **Input**: Samples $\{x_t\}_{t\in\mathcal{S}}$, number of types r, number of iterations T
2: $\tilde{M}_2 \leftarrow \textsc{MatrixAltMin}\left(\hat{M}_2\big(1, \frac{|\mathcal{S}|}{2}\big), r, T\right)$ (see Algorithm 3)
3: Compute eigenvalue decomposition of $\tilde{M}_2 = \tilde{U}_{M_2}\tilde{\Sigma}_{M_2}\tilde{U}_{M_2}^T$
4: $\tilde{H} \leftarrow \textsc{TensorLS}\left(\hat{M}_3\big(\frac{|\mathcal{S}|}{2}+1, |\mathcal{S}|\big), \tilde{U}_{M_2}, \tilde{\Sigma}_{M_2}\right)$ (see Algorithm 4)
5: Compute rank-r decomposition $\sum_{a\in[r]} \hat{\lambda}_a^{\tilde{H}}(\hat{v}_a^{\tilde{H}} \otimes \hat{v}_a^{\tilde{H}} \otimes \hat{v}_a^{\tilde{H}})$ of $\tilde{H}$, using RTPM of [2]
6: **Output**: $\hat{P} = \tilde{U}_{M_2}\tilde{\Sigma}_{M_2}^{1/2}\hat{V}^{\tilde{H}}\hat{\Lambda}^{\tilde{H}}$, $\hat{Q} = (\hat{\Lambda}^{\tilde{H}})^{-2}$, where $\hat{V}^{\tilde{H}} = [\hat{v}_1^{\tilde{H}} \ \dots \ \hat{v}_r^{\tilde{H}}]$ and $\hat{\Lambda}^{\tilde{H}} = \mathrm{diag}(\lambda_1^{\tilde{H}}, \dots, \lambda_r^{\tilde{H}})$

then there exists a permutation π over $[r]$ such that Algorithm 2 achieves the following bounds with a choice of $T = C' \log(N|\mathcal{S}|)$ for all $i \in [r]$, with probability at least $1-\delta$:

$$|\hat{q}_{\pi_i} - q_i| \ \le \ \varepsilon \ , \ \textit{and} \qquad \|\hat{P}_{\pi_i} - P_i\| \ \le \ \varepsilon \sqrt{\frac{r\, q_{\max}\, \sigma_1(M_2)}{q_{\min}}} \ ,$$

where $\mu = \mu(M_2)$ defined in (8) *with run-time* $\mathrm{poly}(N, r, 1/q_{\min}, 1/\varepsilon, \log(1/\delta), \sigma_1(M_2)/\sigma_r(M_2))$.

Algorithm 3 MATRIXALTMIN: Alternating Minimization for Matrix Completion [12]

1: **Input**: $\hat{M}_2\big(1, \frac{|\mathcal{S}|}{2}\big)$, r, T
2: Initialize $N \times r$ dimensional matrix $U_0 \leftarrow$ top-r eigenvectors of $\mathcal{P}_{\Omega_2}(\hat{M}_2\big(1, \frac{|\mathcal{S}|}{2}\big))$
3: **for all** $\tau = 1$ to $T-1$ **do**
4: $\quad \hat{U}_{\tau+1} = \arg\min_U \|\mathcal{P}_{\Omega_2}(\hat{M}_2\big(1, \frac{|\mathcal{S}|}{2}\big)) - \mathcal{P}_{\Omega_2}(UU_\tau^T)\|_F^2$
5: $\quad [U_{\tau+1}R_{\tau+1}] = \mathrm{QR}(\hat{U}_{\tau+1})$ (standard QR decomposition)
6: **end for**
7: **Output**: $\tilde{M}_2 = (\hat{U}_T)(U_{T-1})^T$

Algorithm 4 TENSORLS: Least Squares method for Tensor Estimation [12]

1: **Input**: $\hat{M}_3\big(\frac{|\mathcal{S}|}{2}+1, |\mathcal{S}|\big)$, $\hat{U}_{M_2}$, $\hat{\Sigma}_{M_2}$
2: Define operator $\hat{\nu} : \mathbb{R}^{r\times r\times r} \to \mathbb{R}^{N\times N\times N}$ as follows

$$\hat{\nu}_{ijk}(Z) = \begin{cases} \sum_{abc} Z_{abc}(\hat{U}_{M_2}\hat{\Sigma}_{M_2}^{1/2})_{ia}(\hat{U}_{M_2}\hat{\Sigma}_{M_2}^{1/2})_{jb}(\hat{U}_{M_2}\hat{\Sigma}_{M_2}^{1/2})_{kc}, & \text{if } i\neq j\neq k\neq i \ , \\ 0, & \text{otherwise.} \end{cases} \tag{13}$$

3: Define $\hat{A} : \mathbb{R}^{r\times r\times r} \to \mathbb{R}^{r\times r\times r}$ s.t. $\hat{A}(Z) = \hat{\nu}(Z)[\hat{U}_{M_2}\hat{\Sigma}_{M_2}^{-1/2}, \hat{U}_{M_2}\hat{\Sigma}_{M_2}^{-1/2}, \hat{U}_{M_2}\hat{\Sigma}_{M_2}^{-1/2}]$
4: **Output**: $\arg\min_Z \|\hat{A}(Z) - \mathcal{P}_{\Omega_3}\big(\hat{M}_3\big(\frac{|\mathcal{S}|}{2}+1, |\mathcal{S}|\big)\big)[\hat{U}_{M_2}\hat{\Sigma}_{M_2}^{-1/2}, \hat{U}_{M_2}\hat{\Sigma}_{M_2}^{-1/2}, \hat{U}_{M_2}\hat{\Sigma}_{M_2}^{-1/2}]\|_F^2$

3.2 Phase 2: RANKCENTRALITY.

Recall that $E = \{(i_k, j_k) : i_k \neq j_k \in [n], 1 \le k \le N\}$ represents collection of $N = |E|$ pairs and $G = ([n], E)$ is the corresponding graph. Let $\tilde{P}_a$ denote the estimation of $P_a = [P_{ka}] \in [-1, 1]^N$ for the mixture component $a, 1 \le a \le r$; where P_{ka} is defined as per (1). For each a, using G and $\tilde{P}_a$, we shall use the RANKCENTRALITY [19] to obtain estimation of $\mathbf{w}^{(a)}$. Next we describe the algorithm and guarantees associated with it.

Without loss of generality, we can assume that $\mathbf{w}^{(a)}$ is such that $\sum_i w_i^{(a)} = 1$ for all $a, 1 \le a \le r$. Given this normalization, RANKCENTRALITY estimates $\mathbf{w}^{(a)}$ as stationary distribution of an appropriate Markov chain on G. The transition probabilities are 0 for all $(i,j) \notin E$. For $(i,j) \in E$, they are function of $\tilde{P}_a$. Specifically, transition matrix $\tilde{p}^{(a)} = [\tilde{p}_{i,j}^{(a)}] \in [0,1]^{n\times n}$ with $\tilde{p}_{i,j}^{(a)} = 0$ if

$(i,j) \notin E$, and for $(i_k, j_k) \in E$ for $1 \le k \le N$,

$$\tilde{p}^{(a)}_{i_k,j_k} = \frac{1}{d_{\max}} \frac{(1+\tilde{P}_{ka})}{2} \quad \text{and} \quad \tilde{p}^{(a)}_{j_k,i_k} = \frac{1}{d_{\max}} \frac{(1-\tilde{P}_{ka})}{2}, \tag{14}$$

Finally, $\tilde{p}^{(a)}_{i,i} = 1 - \sum_{j \neq i} \tilde{p}^{(a)}_{i,j}$ for all $i \in [n]$. Let $\tilde{\pi}^{(a)} = [\tilde{\pi}^{(a)}_i]$ be a stationary distribution of the Markov chain defined by $\tilde{p}^{(a)}$. That is,

$$\tilde{\pi}^{(a)}_i = \sum_j \tilde{p}^{(a)}_{ji}\, \tilde{\pi}^{(a)}_j \quad \text{for all } i \in [n]. \tag{15}$$

Computationally, we suggest obtaining estimation of $\tilde{\pi}$ by using power-iteration for T iterations. As argued before, cf. [19], $T = \Theta\big(b^2 d_{\max}(\log n + \log(1/\varepsilon))/(\xi d_{\min})\big)$, is sufficient to obtain reasonably good estimation of $\tilde{\pi}$.

The underlying assumption here is that there is a unique stationary distribution, which is established by our result under the conditions of Theorem 1. Now $\tilde{p}$ is an approximation of the ideal transition probabilities, where $p^{(a)} = [p^{(a)}_{i,j}]$ where $p^{(a)}_{i,j} = 0$ if $(i,j) \notin E$ and $p^{(a)}_{i,j} \propto w^{(a)}_j/(w^{(a)}_i + w^{(a)}_j)$ for all $(i,j) \in E$. Such an ideal Markov chain is *reversible* and as long as G is connected (which is, in our case, by choice), the stationary distribution of this ideal chain is $\pi^{(a)} = \mathbf{w}^{(a)}$ (recall, we have assumed $\mathbf{w}^{(a)}$ to be normalized so that all its components up to 1).

Now $\tilde{p}^{(a)}$ is an approximation of such an ideal transition matrix $p^{(a)}$. In what follows, we state result about how this approximation error translates into the error between $\tilde{\pi}^{(a)}$ and $\mathbf{w}^{(a)}$. Recall that $b \equiv \max_{i,j \in [n]} w_i/w_j$, $d_{\max}$ and $d_{\min}$ are maximum and minimum vertex degrees of G and ξ as defined in (9).

Theorem 4. *Let $G = ([n], E)$ be non-bipartite and connected. Let $\|\tilde{p}^{(a)} - p^{(a)}\|_2 \le \varepsilon$ for some positive $\varepsilon \le (1/4)\xi b^{-5/2}(d_{\min}/d_{\max})$. Then, for some positive universal constant C,*

$$\frac{\|\tilde{\pi}^{(a)} - \mathbf{w}^{(a)}\|}{\|\mathbf{w}^{(a)}\|} \le \frac{C\, b^{5/2}}{\xi} \frac{d_{\max}}{d_{\min}}\, \varepsilon. \tag{16}$$

And, starting from any initial condition, the power iteration manages to produce an estimate of $\tilde{\pi}^{(a)}$ within twice the above stated error bound in $T = \Theta\big(b^2 d_{\max}(\log n + \log(1/\varepsilon))/(\xi d_{\min})\big)$ iterations.

Proof of the above result can be found in Appendix. For spectral expander (e.g. connected Erdos-Renyi graph with high probability), $\xi = \Theta(1)$ and therefore the bound is effectively $O(\varepsilon)$ for bounded dynamic range, i.e. $b = O(1)$.

4 Discussion

Learning distribution over permutations of n objects from partial observation is fundamental to many domains. In this work, we have advanced understanding of this question by characterizing sufficient conditions and associated algorithm under which it is feasible to learn mixed MNL model in computationally and statistically efficient (polynomial in problem size) manner from partial/pairwise comparisons. The conditions are natural – the mixture components should be "identifiable" given partial preference/comparison data – stated in terms of full rank and incoherence conditions of the second moment matrix. The algorithm allows learning of mixture components as long as number of mixture components scale $o(n^{2/7})$ for distribution over permutations of n objects.

To the best of our knowledge, this work provides first such sufficient condition for learning mixed MNL model – a problem that has remained open in econometrics and statistics for a while, and more recently Machine learning. Our work nicely complements the impossibility results of [1].

Analytically, our work advances the recently popularized spectral/tensor approach for learning mixture model from lower order moments. Concretely, we provide means to learn the component even when only partial information about the sample is available unlike the prior works. To learn the model parameters, once we identify the moments associated with each mixture, we advance the result of [19] in its applicability. Spectral methods have also been applied to ranking in the context of assortment optimization in [5].

References

[1] A. Ammar, S. Oh, D. Shah, and L. Voloch. What's your choice? learning the mixed multi-nomial logit model. In *Proceedings of the ACM SIGMETRICS/international conference on Measurement and modeling of computer systems*, 2014.

[2] A. Anandkumar, R. Ge, D. Hsu, S. M. Kakade, and M. Telgarsky. Tensor decompositions for learning latent variable models. *CoRR*, abs/1210.7559, 2012.

[3] H. Azari Soufiani, W. Chen, D. C Parkes, and L. Xia. Generalized method-of-moments for rank aggregation. In *Advances in Neural Information Processing Systems 26*, pages 2706–2714. 2013.

[4] H. Azari Soufiani, D. Parkes, and L. Xia. Computing parametric ranking models via rank-breaking. In *Proceedings of The 31st International Conference on Machine Learning*, pages 360–368, 2014.

[5] J. Blanchet, G. Gallego, and V. Goyal. A markov chain approximation to choice modeling. In *EC*, pages 103–104, 2013.

[6] B. Bollobás. *Random Graphs*. Cambridge University Press, January 2001.

[7] R. A. Bradley and M. E. Terry. Rank analysis of incomplete block designs: I. the method of paired comparisons. *Biometrika*, 39(3/4):324–345, 1955.

[8] E. J. Candès and B. Recht. Exact matrix completion via convex optimization. *Foundations of Computational Mathematics*, 9(6):717–772, 2009.

[9] C. Davis and W. M. Kahan. The rotation of eigenvectors by a perturbation. iii. *SIAM Journal on Numerical Analysis*, 7(1):1–46, 1970.

[10] J. C. Duchi, L. Mackey, and M. I. Jordan. On the consistency of ranking algorithms. In *Proceedings of the ICML Conference*, Haifa, Israel, June 2010.

[11] V. F. Farias, S. Jagabathula, and D. Shah. A data-driven approach to modeling choice. In *NIPS*, pages 504–512, 2009.

[12] P. Jain and S. Oh. Learning mixtures of discrete product distributions using spectral decompositions. *arXiv preprint arXiv:1311.2972*, 2014.

[13] L. R. Ford Jr. Solution of a ranking problem from binary comparisons. *The American Mathematical Monthly*, 64(8):28–33, 1957.

[14] R. H. Keshavan, A. Montanari, and S. Oh. Matrix completion from a few entries. *Information Theory, IEEE Transactions on*, 56(6):2980–2998, 2010.

[15] R. H. Keshavan, A. Montanari, and S. Oh. Matrix completion from noisy entries. *The Journal of Machine Learning Research*, 99:2057–2078, 2010.

[16] D. R. Luce. *Individual Choice Behavior*. Wiley, New York, 1959.

[17] D. McFadden. Conditional logit analysis of qualitative choice behavior. *Frontiers in Econometrics*, pages 105–142, 1973.

[18] I. Mitliagkas, A. Gopalan, C. Caramanis, and S. Vishwanath. User rankings from comparisons: Learning permutations in high dimensions. In *Communication, Control, and Computing (Allerton), 2011 49th Annual Allerton Conference on*, pages 1143–1150. IEEE, 2011.

[19] S. Negahban, S. Oh, and D. Shah. Iterative ranking from pair-wise comparisons. In *NIPS*, pages 2483–2491, 2012.

[20] S. Negahban and M. J. Wainwright. Restricted strong convexity and (weighted) matrix completion: Optimal bounds with noise. *Journal of Machine Learning Research*, 2012.

[21] P. Samuelson. A note on the pure theory of consumers' behaviour. *Economica*, 5(17):61–71, 1938.

[22] H. A. Soufiani, D. C. Parkes, and L. Xia. Random utility theory for social choice. In *NIPS*, pages 126–134, 2012.

[23] Louis L Thurstone. A law of comparative judgment. *Psychological review*, 34(4):273, 1927.

[24] J. Tropp. User-friendly tail bounds for sums of random matrices. *Foundations of Computational Mathematics*, 2011.

[25] E. Zermelo. Die berechnung der turnier-ergebnisse als ein maximumproblem der wahrscheinlichkeitsrechnung. *Mathematische Zeitschrift*, 29(1):436–460, 1929.

Supplementary material for "Learning Mixed Multinomial Logit Model from Ordinal Data"

A Proof of Theorem 1

In order to apply Theorem 4, let $\Delta = p - \tilde{p}$, then

$$\begin{aligned}\|\Delta\|_2 \leq \|\Delta - \mathrm{diag}(\Delta)\|_2 \ + \ \|\mathrm{diag}(\Delta)\|_2 \ &\leq \|\Delta - \mathrm{diag}(\Delta)\|_F \ + \ \max_{i\in[n]} \Delta_{ii} \\ &\leq \sqrt{\sum_{i\neq j}(p_{ij} - \tilde{p}_{ij})^2} \ + \ \max_{i\in[n]} \Big| \sum_{j\neq i}(p_{ij} - \tilde{p}_{ij}) \Big| \\ &\leq \frac{1}{\sqrt{2}\, d_{\max}} \|P_a - \tilde{P}_a\| \ + \ \frac{1}{2\sqrt{d_{\max}}} \|P_a - \tilde{P}_a\| \ .\end{aligned}$$

From Theorem 3, we know that $\|\Delta\|_2 \leq \frac{2}{\sqrt{d_{\max}}} \|P_a - \tilde{P}_a\| \leq \varepsilon \sqrt{(r q_{\max} \sigma_1(M_2))/q_{\min}}$ and substituting this into the bound in Theorem 4 we get the desired bound.

B Proof of the performance guarantee for the spectral method in Theorem 3

To simplify notations, we will assume that the indices of the output of the algorithm and the ground truths are matched such that the theorem holds with identity permutation. The spectral algorithm outputs $\hat{P} = \tilde{U}_{M_2} \tilde{\Sigma}_{M_2}^{1/2} \hat{V}^{\tilde{H}} \hat{\Lambda}^{\tilde{H}}$. From theorem 2, we know that $P = U_{M_2} \Sigma_{M_2}^{1/2} V^H \Lambda^H$, or equivalently $P = U_{M_2} \Sigma_{M_2}^{1/2} V^H Q^{-1/2}$. To show that P and $\hat{P}$ are close, we would hope that each of the terms above to be close.

To that end, define

$$\widetilde{V} \ \equiv \ \tilde{\Sigma}_{M_2}^{-1/2} \tilde{U}_{M_2}^T P Q^{1/2} \ , \tag{17}$$

$$\widetilde{G} \ \equiv \ \sum_{i=1}^{r} \frac{1}{\sqrt{q_i}} (\tilde{v}_i \otimes \tilde{v}_i \otimes \tilde{v}_i) \ , \tag{18}$$

where $\widetilde{V} = [\tilde{v}_1 \, \tilde{v}_2 \, \ldots \tilde{v}_r]$ (note that $\widetilde{G}$ is proxy of V^H). Now

$$\begin{aligned}\big\| \hat{P} - P \big\|_2 \ &\leq \ \big\| \tilde{U}_{M_2} \tilde{U}_{M_2}^T P - P \big\|_2 + \big\| \hat{P} - \tilde{U}_{M_2} \tilde{U}_{M_2}^T P \big\|_2 \\ &= \ \big\| \tilde{U}_{M_2} \tilde{U}_{M_2}^T P - P \big\|_2 + \big\| \tilde{U}_{M_2} \tilde{\Sigma}_{M_2}^{1/2} \hat{V}^{\tilde{H}} \hat{\Lambda}^{\tilde{H}} - \tilde{U}_{M_2} \tilde{\Sigma}_{M_2}^{1/2} \widetilde{V} Q^{-1/2} \big\|_2 \qquad \text{using (17)} \\ &\leq \ \big\| \tilde{U}_{M_2} \tilde{U}_{M_2}^T P - P \big\|_2 + \big\| \tilde{U}_{M_2} \tilde{\Sigma}_{M_2}^{1/2} (\widetilde{V} - \hat{V}^{\tilde{H}}) Q^{-1/2} \big\|_2 \\ &\qquad + \big\| \tilde{U}_{M_2} \tilde{\Sigma}_{M_2}^{1/2} \hat{V}^{\tilde{H}} (Q^{-1/2} - \hat{\Lambda}^{\tilde{H}}) \big\|_2 \ . \end{aligned} \tag{19}$$

To bound the three terms on the RHS of (19), we shall use the following 'errors' (which we shall bound sharply later in the proof): define (recall $\tilde{H}$ was produced by Algorithm Tensor Least Squares)

$$\begin{aligned}\varepsilon_M &= \frac{\|M_2 - \tilde{M}_2\|_2}{\sigma_r(M_2)} \\ \varepsilon_H &= \|\tilde{H} - H\|_2 . \end{aligned} \tag{20}$$

Bounding the first term in RHS of (19)**.** To begin with, note that we can represent $P = U_{M_2} \Sigma V^T$ by definition. Given the definition of (20), and an application of Davis-Kahan theorem [9] implies that

$$\|(\tilde{U}_{M_2} \tilde{U}_{M_2}^T - \mathbb{I}) U_{M_2}\|_2 \leq \varepsilon_M . \tag{21}$$

Using this, we have

$$\begin{aligned}\|\tilde{U}_{M_2} \tilde{U}_{M_2}^T P - P\|_2 &= \|(\tilde{U}_{M_2} \tilde{U}_{M_2}^T - \mathbb{I}) P\|_2 \\ &\leq \|(\tilde{U}_{M_2} \tilde{U}_{M_2}^T - \mathbb{I})\|_2 U_{M_2} \|\Sigma V^T\|_2 \\ &\leq \varepsilon_M \|P\|_2 \ , \end{aligned} \tag{22}$$

Bounding the second term in RHS of (19)**.** Consider

$$\begin{aligned}\|\tilde{\Sigma}_{M_2}\|_2 = \|\tilde{M}_2\|_2 &\leq \|\tilde{M}_2 - M_2\|_2 + \|M_2\|_2 \\ &\leq \varepsilon_M \sigma_r(M_2) + \|M_2\|_2. \end{aligned} \tag{23}$$

We state the following Lemma providing bound on $\|\widetilde{V} - \hat{V}^{\tilde{H}}\|_2$ in terms of ε_M and ε_G:

Lemma B.1. *There exists universal constant $C_1 > 0$ such that*

$$\begin{aligned}\|\widetilde{V} - \hat{V}^{\tilde{H}}\|_2 &\leq C_1\sqrt{r\, q_{\max}}\left(\varepsilon_H + \frac{1}{\sqrt{q_{\min}}}\,\varepsilon_M\right), \text{ and} \\ \|Q^{-1/2} - \hat{\Lambda}^H\|_2 &\leq C_1\Big(\varepsilon_H + \frac{1}{\sqrt{q_{\min}}}\,\varepsilon_M\Big).\end{aligned}$$

Given (23), Lemma B.1, the fact that $\tilde{U}_{M_2}$ is a unitary matrix and Q is a diagonal matrix, we obtain that the second term in RHS of (19) is bounded by

$$\begin{aligned}\big\|\,\tilde{U}_{M_2}\tilde{\Sigma}_{M_2}^{1/2}\,(\widetilde{V} - \hat{V}^{\tilde{H}})\,Q^{-1/2}\big\|_2 &\leq \big\|\,\tilde{U}_{M_2}\big\|_2\big\|\tilde{\Sigma}_{M_2}^{1/2}\,\big\|_2\big\|(\widetilde{V} - \hat{V}^{\tilde{H}})\,\big\|_2\big\|Q^{-1/2}\big\|_2 \\ &\leq \frac{\sqrt{\varepsilon_M\sigma_r(M_2) + \|M_2\|_2}}{\sqrt{q_{\min}}}\,\|\widetilde{V} - \hat{V}^{\tilde{H}}\|_2 \\ &\leq C_2\sqrt{\frac{\|M_2\|_2\, r\, q_{\max}}{q_{\min}}}\Big(\varepsilon_H + \frac{1}{\sqrt{q_{\min}}}\,\varepsilon_M\Big), \end{aligned} \tag{24}$$

for an appropriate universal constant $C_2 > 0$, with $\varepsilon_M \leq 1/2$ and using fact $\sigma_r(M_2) \leq \|M_2\|_2$.

Bounding the third term in RHS of (19)**.** Observe that

$$\begin{aligned}\|\hat{V}^{\tilde{H}}\|_2 &\leq \|\widetilde{V}\|_2 + \|\hat{V}^{\tilde{H}} - \widetilde{V}\|_2 \\ &\leq 1 + 2\varepsilon_M + C_1\sqrt{r q_{\max}}(\varepsilon_H + \varepsilon_M/\sqrt{q_{\min}}) \end{aligned} \tag{25}$$

using Remark C.1 and Lemma B.1. Using (25) and Lemma B.1, the last term in (19) can be bounded above as

$$\begin{aligned}\big\|\tilde{U}_{M_2}\tilde{\Sigma}_{M_2}^{1/2}\hat{V}^{\tilde{H}}\,(Q^{-1/2} - \hat{\Lambda}^{\tilde{H}})\,\big\|_2 &\leq \big\|\tilde{U}_{M_2}\tilde{\Sigma}_{M_2}^{1/2}\big\|_2\big\|\hat{V}^{\tilde{H}}\,\big\|_2\big\|(Q^{-1/2} - \hat{\Lambda}^{\tilde{H}})\,\big\|_2 \\ &\leq \left(\sqrt{\varepsilon_M\sigma_r(M_2) + \|M_2\|_2}\right)\|\hat{V}^{\tilde{H}}\|_2\,\|Q^{-1/2} - \hat{\Lambda}^G\|_2 \\ &\leq C_3\sqrt{\|M_2\|_2}\left(1 + \sqrt{r q_{\max}}\varepsilon_H + \frac{\sqrt{r q_{\max}}}{\sqrt{q_{\min}}}\,\varepsilon_M\right)\left(\varepsilon_H + \frac{1}{\sqrt{q_{\min}}}\,\varepsilon_M\right) \qquad (26) \\ &\leq C_3\sqrt{\|M_2\|_2}\left(\varepsilon_H + \frac{1}{\sqrt{q_{\min}}}\,\varepsilon_M\right), \qquad (27)\end{aligned}$$

for $\varepsilon_M \leq \sqrt{q_{\min}/(r\, q_{\max})}$, $\varepsilon_H \leq 1/\sqrt{r\, q_{\max}}$ and for some universal constant $C_3 > 0$.

Towards Theorem 3. Substituting (22), (24), (27) in (19), for universal constant $C_4 > 0$, we get

$$\begin{aligned}\|\hat{P} - P\|_2 &\leq \varepsilon_M\|P\|_2 + C_4\sqrt{\frac{\|M_2\|_2\, r\, q_{\max}}{q_{\min}}}\Big(\varepsilon_H + \frac{1}{\sqrt{q_{\min}}}\,\varepsilon_M\Big) \\ &\leq C_4\sqrt{\frac{\|M_2\|_2\, r\, q_{\max}}{q_{\min}}}\Big(\varepsilon_H + \frac{1}{\sqrt{q_{\min}}}\,\varepsilon_M\Big), \end{aligned} \tag{28}$$

where we used the fact that $\|P\|_2 \leq \|\sqrt{\|M_2\|_2/q_{\min}}\|$ since $P = U_{M_2}\Sigma_{M_2}^{1/2}V^H\Lambda^H$. Given (28), to complete the proof of Theorem 3, we need to establish bounds on ε_M and ε_H.

Bounding ε_M. To bound ε_M, we need to bound error between M_2 and $\tilde{M}_2$, the output of alternating minimization procedure applied to $\hat{M}_2$. The following theorem [12] provides such a bound.

Theorem 5 (Theorem 4.1, [12])**.** *For an $N \times N$ symmetric rank-r matrix M with incoherence μ, we observe off-diagonal entries corrupted by noise:*

$$\hat{M}_{ij} = \begin{cases} M_{ij} + E_{ij} & \text{if } i \neq j\,, \\ 0 & \text{otherwise.}\end{cases}$$

Let $\hat{M}^{(\tau)}$ denote the output after τ iterations of MATRIXALTMIN. *If $\mu \leq (\sigma_r(M)/\sigma_1(M))\sqrt{N/(32\, r^{1.5})}$, the noise is bounded by $\|\mathcal{P}_{\Omega_2}(E)\|_2 \leq \sigma_r(M)/32\sqrt{r}$, and each column of the noise is bounded by $\|\mathcal{P}_{\Omega_2}(E)_i\| \leq \sigma_1(M)\mu\sqrt{3r/(8\,N)}$, $\forall i \in [N]$, then after $\tau \geq (1/2)\log\big(2\|M\|_F/\varepsilon\big)$ iterations of* MATRIXALTMIN, *the estimate $\hat{M}^{(\tau)}$ satisfies:*

$$\|M - \hat{M}^{(\tau)}\|_2 \quad \leq \quad \varepsilon + \frac{9\,\|M\|_F\,\sqrt{r}}{\sigma_r(M)}\|\mathcal{P}_{\Omega_2}(E)\|_2\ ,$$

for any $\varepsilon \in (0,1)$. Further, $\hat{M}^{(\tau)}$ is μ_1-incoherent with $\mu_1 \leq 6\mu\sigma_1(M_2)/\sigma_r(M_2)$.

To apply the above result in our setting to bound ε_M, we need to bound $\|\mathcal{P}_{\Omega_2}(E)\|_2$ and $\|\mathcal{P}_{\Omega_2}(E)_i\|$ for all $i \in [N]$. To that end, we state the following Lemma.

Lemma B.2. *Let $S_2 \equiv (2/|\mathcal{S}|)\sum_{t\in\{1,\ldots,|\mathcal{S}|/2\}} x_t x_t^T$ be the sample covariance matrix, and let $E = \frac{N(N-1)}{\ell(\ell-1)}S_2 - M_2$ denote the sampling error in the off-diagonal entries. Then, there exists a universal constant $C_5 > 0$ such that with probability at least $1-\delta$,*

$$\|\mathcal{P}_{\Omega_2}(E)\|_2 \quad \leq \quad C_5\sqrt{\frac{N^2\log(N/\delta)}{\ell|\mathcal{S}|}\Big(\sigma_1(M_2) + \frac{N}{\ell}\Big)}\ .$$

Moreover, the Euclidean norm of the columns are uniformly bounded by

$$\|\mathcal{P}_{\Omega_2}(E)_i\| \quad \leq \quad C_5\sqrt{\frac{N^3\log(N/\delta)}{\ell^2|\mathcal{S}|}}\ ,$$

for all $i \in [N]$.

Theorem 5 and Lemma B.2 imply that with probability at least $1-\delta$, for large enough iterations of the MATRIXALTMIN,

$$\varepsilon_M = \frac{1}{\sigma_r(M_2)}\|M_2 - \tilde{M}_2\|_2 \ \leq\ C_6\frac{\|M_2\|_F\,N}{\sigma_r(M_2)^2}\sqrt{\frac{r\ \log(N/\delta)}{\ell\,|\mathcal{S}|}\Big(\sigma_1(M_2) + \frac{N}{\ell}\Big)}\ , \tag{29}$$

for universal constant $C_6, C_7 > 0$ when $|\mathcal{S}| \geq C_7 N^4 r\log(N/\delta)/(\ell^2\sigma_r(M_2)^2)$ – the assumption of Theorem statement.

Bounding ε_H. To bound ε_H, we need bound on error induced by the output of the Tensor Least Square procedure. The following result [12] provides such a bound.

Theorem 6 (Theorem 4.3, [12])**.** *If $N \geq \frac{144 r^3 \sigma_1(M_2)^2}{\sigma_r(M_2)^2}$, then with probability at least $1-\delta$,*

$$\|H - \tilde{H}\|_F \quad \leq \quad \frac{24\mu_1^3\mu r^{3.5}\sigma_1(M_2)^{3/2}}{N q_{\min}^{1/2}\sigma_r(M_2)^{3/2}}\,\varepsilon_M + 2\Big\|\,\mathcal{P}_{\Omega_3}\big(S_3 - M_3\big)\big[\tilde{U}_{M_2}\tilde{\Sigma}_{M_2}^{-1/2}, \tilde{U}_{M_2}\tilde{\Sigma}_{M_2}^{-1/2}, \tilde{U}_{M_2}\tilde{\Sigma}_{M_2}^{-1/2}\big]\,\Big\|_F\ ,$$

where $\varepsilon_M = (1/\sigma_r(M_2))\|\tilde{M}_2 - M_2\|$, $\mu = \mu(M_2)$, $\mu_1 = \mu(\tilde{U}_{M_2})$, and

$$S_3 = \frac{N(N-1)(N-2)}{\ell(\ell-1)(\ell-2)}\frac{2}{|\mathcal{S}|}\Big(\sum_{t=1+|\mathcal{S}|/2}^{|\mathcal{S}|} x_t\otimes x_t\otimes x_t\Big). \tag{30}$$

To utilize above result, we state the following Lemma.

Lemma B.3. *There exists a positive numerical constant C_9 such that with probability at least $1-\delta$,*

$$\Big\|\,\mathcal{P}_{\Omega_3}\big(S_3 - M_3\big)\big[\tilde{U}_{M_2}\tilde{\Sigma}_{M_2}^{-1/2}, \tilde{U}_{M_2}\tilde{\Sigma}_{M_2}^{-1/2}, \tilde{U}_{M_2}\tilde{\Sigma}_{M_2}^{-1/2}\big]\,\Big\|_F \quad \leq \quad C_9\frac{r^3\mu_1^3 N^{3/2}}{\sigma_r(M_2)^{3/2}}\sqrt{\frac{\log(1/\delta)}{|\mathcal{S}|}}\ .$$

Theorem 6 and Lemma B.3 imply that with probability at least $1-\delta$

$$\begin{aligned}\varepsilon_H &= \|H - \tilde{H}\|_2 \\ &\leq C_{10}\,\mu_1^3\,r^3\,\left(\frac{\mu r^{1/2}}{N q_{\min}^{1/2}}\frac{\sigma_1(M_2)^{3/2}}{\sigma_r(M_2)^{3/2}}\varepsilon_M + \frac{N^{3/2}}{\sigma_r(M_2)^{3/2}}\sqrt{\frac{\log(1/\delta)}{|\mathcal{S}|}}\right)\end{aligned} \tag{31}$$

for some positive numerical constants C_{10}, C_{11} when $N \geq C_{11} \frac{r^3 \sigma_1(M_2)^2}{\sigma_r(M_2)^2}$.

Equations (28)-(31) imply that

$$\|\hat{P} - P\|_2 \leq C_4 \sqrt{\frac{\|M_2\|_2 \, r \, q_{\max}}{q_{\min}}} \Big(\varepsilon_H + \frac{1}{\sqrt{q_{\min}}} \, \varepsilon_M \Big) \, ,$$

$$\leq C_4 \sqrt{\frac{\|M_2\|_2 \, r \, q_{\max}}{q_{\min}}} \Big(C_{10} \, \mu_1^3 \, r^3 \, \left(\frac{\mu r^{1/2}}{N q_{\min}^{1/2}} \frac{\sigma_1(M_2)^{3/2}}{\sigma_r(M_2)^{3/2}} \varepsilon_M + \frac{N^{3/2}}{\sigma_r(M_2)^{3/2}} \sqrt{\frac{\log(1/\delta)}{|\mathcal{S}|}} \right) + \frac{1}{\sqrt{q_{\min}}} \, \varepsilon_M \Big) \, , .$$

Using $\mu_1 \leq 6\mu\sigma_1(M_2)/\sigma_r(M_2)$ from Theorem 5, we obtain (for appropriate constant $C_{12} > 0$),

$$\|\hat{P} - P\|_2 \leq C_{12} \sqrt{\frac{\|M_2\|_2 \, r \, q_{\max}}{q_{\min}^2}} \left\{ \Big(\frac{\mu^4 r^{3.5}}{N} \Big(\frac{\sigma_1(M_2)}{\sigma_r(M_2)} \Big)^{4.5} + 1 \Big) \varepsilon_M + \frac{\mu^3 r^3 \sigma_1(M_2)^3 N^{3/2} q_{\min}^{1/2}}{\sigma_r(M_2)^{4.5}} \sqrt{\frac{\log(1/\delta)}{|\mathcal{S}|}} \right\}.$$

$$\leq C_{12} \sqrt{\frac{\|M_2\|_2 \, r \, q_{\max}}{q_{\min}}} \left\{ \frac{2\varepsilon_M}{q_{\min}^{1/2}} + \mu^3 r^3 \frac{\sigma_1(M_2)^3 N^{1.5}}{\sigma_r(M_2)^{4.5}} \sqrt{\frac{\log(1/\delta)}{|\mathcal{S}|}} \right\} \, , \tag{32}$$

where $N \geq \mu^4 r^{3.5} (\sigma_1(M_2)/\sigma_r(M_2))^{4.5}$ as per assumption of the Theorem statement. From (29), it follows that when

$$|\mathcal{S}| \geq C_{13} \frac{r^2 \sigma_1(M_2)^2 N^2 (\sigma_1(M_2) + N/\ell) \log(N/\delta)}{\varepsilon^2 q_{\min} \ell \sigma_r(M_2)^4} \geq C_{13} \frac{r \|M_2\|_F^2 N^2 \log(N/\delta)}{\varepsilon^2 q_{\min} \ell \sigma_r(M_2)^4} \Big(\sigma_1(M_2) + \frac{N}{\ell} \Big), \tag{33}$$

(we used $\|M_2\|_F \leq \sqrt{r}\sigma_1(M_2)$ for rank r matrix M_2) for an appropriate choice of universal constant $C_{13} > 0$,

$$\frac{2\varepsilon_M}{q_{\min}^{1/2}} \leq \frac{\varepsilon}{2}. \tag{34}$$

Also, recall that for (29) to hold, we require

$$|\mathcal{S}| \geq C_7 \frac{N^4 r \log(N/\delta)}{\ell^2 \sigma_r(M_2)^2} \, . \tag{35}$$

Also, when

$$|\mathcal{S}| \geq C_{14} \frac{\mu^6 r^6 N^3 \sigma_1(M_2)^6 \log(1/\delta)}{\varepsilon^2 \sigma_r(M_2)^9}, \tag{36}$$

for an appropriate choice of universal constant $C_{14} > 0$, the second term inside bracket in (32) is less than $\frac{\varepsilon}{2}$. From above, it follows that when

$$|\mathcal{S}| \geq C_{15} \frac{r^2 N^2 \sigma_1(M_2)^2 \log(N/\delta)}{q_{\min} \varepsilon^2 \sigma_r(M_2)^4} \Big(\frac{N}{\ell^2} + \frac{\sigma_1(M_2)}{\ell} + \frac{q_{\min} \mu^6 r^4 \sigma_1(M_2)^4 N}{\sigma_r(M_2)^5} \Big) + \frac{N^4 r \log(N/\delta)}{\ell^2 \sigma_1(M_2)^2} \tag{37}$$

$$\begin{aligned} \|\hat{P} - P\|_2 &\leq \varepsilon \sqrt{\frac{r \, q_{\max} \sigma_1(M_2)}{q_{\min}}} \, , \text{ and} \\ \|Q^{-1/2} - \hat{\Lambda}^G\|_2 &\leq \varepsilon \, , \end{aligned}$$

for any $\varepsilon \in (0, C')$ for some positive constants $C = C_{15}$ and C'. Assuming $N \geq C'' r^{3.5} \mu^6 (\sigma_1(M_2)/\sigma_r(M_2))^{4.5}$, the above holds for

$$|\mathcal{S}| \geq C''' \frac{r N^4 \log(N/\delta)}{q_{\min} \sigma_1(M_2)^2 \varepsilon^2} \Big(\frac{1}{\ell^2} + \frac{\sigma_1(M_2)}{\ell N} + \frac{r^4 \sigma_1(M_2)^4}{\sigma_r(M_2)^5} \Big) \, .$$

C Proof of the technical lemmas for the spectral method

C.1 Proof of Lemma B.1

In order to apply the perturbation analysis of Theorem C.2 from [2], it is crucial that we compare to a tensor with an orthogonal decomposition. Since both $\tilde{H}$ and $\widetilde{G}$ do not have orthogonal decompositions, we define a new tensor $\bar{G}$ that is close to $\widetilde{G}$ and has an orthogonal decomposition. Given the singular value decomposition of $\widetilde{V} = XSY^T$, define $\bar{V} \equiv XY^T$. This $\bar{V} \in \mathbb{R}^{r\times r}$ is orthogonal such that $\bar{V}\bar{V}^T = \bar{V}^T\bar{V} = \mathbb{I}$, and is close to $\widetilde{V}$ such that

$$\begin{aligned} \|\widetilde{V} - \bar{V}\|_2 &= \|X(S-\mathbb{I})Y^T\|_2 \\ &\leq \max_{i\in[r]} |S_{ii} - 1| \\ &\leq 2\varepsilon_M \ , \end{aligned} \tag{38}$$

where the last inequality follows from the next remark.

Remark C.1 (Remark 10 in [12]). *Suppose $\|M_2 - \tilde{M}_2\|_2 \leq \varepsilon_M \sigma_r(M_2)$, then*

$$\|\mathbb{I} - \widetilde{V}\widetilde{V}^T\|_2 \ \leq \ 2\varepsilon_M \ .$$

It follows that $\|S^2 - \mathbb{I}\|_2 = \|\widetilde{V}\widetilde{V}^T - \mathbb{I}\|_2 \leq 2\varepsilon_M$. Therefore, $S_{ii}^2 \in [1-2\varepsilon_M, 1+2\varepsilon_M]$ and so is $S_{ii} \in [1-2\varepsilon_M, 1+2\varepsilon_M]$ for $\varepsilon_M \leq 1/2$.

Since $\|\widetilde{V} - \hat{V}^{\tilde{H}}\|_2 \leq \|\widetilde{V} - \bar{V}\|_2 + \|\bar{V} - \hat{V}^{\tilde{H}}\|_2 \leq 2\varepsilon_M + \|\bar{V} - \hat{V}^{\tilde{H}}\|_2$, we are left to show that $\|\bar{V} - \hat{V}^{\tilde{H}}\|_2 \leq 8\sqrt{r\, q_{\max}}(\varepsilon_H + (13/\sqrt{q_{\min}})\varepsilon_M)$ to finish the proof. Recall that $\widetilde{G} = \sum_{i=1}^r \frac{1}{\sqrt{q_i}}(\tilde{v}_i \otimes \tilde{v}_i \otimes \tilde{v}_i)$ and that $\hat{V}^{\tilde{H}}$ is the output of the robust power method applied to $\tilde{H}$, and let

$$\bar{G} \ \equiv \ \sum_{i=1}^r \frac{1}{\sqrt{q_i}}(\bar{v}_i \otimes \bar{v}_i \otimes \bar{v}_i) \ .$$

Applying (38), we get that

$$\begin{aligned} \|\widetilde{G} - \bar{G}\|_2 &= \max_{\|u\|=1} \Big\| \sum_{i=1}^r \frac{1}{\sqrt{q_i}}(\tilde{v}_i \otimes \tilde{v}_i \otimes \tilde{v}_i - \bar{v}_i \otimes \bar{v}_i \otimes \bar{v}_i) \Big\|_2 \\ &\leq \max_{\|u\|=1} \sum_{i=1}^r \frac{1}{\sqrt{q_i}} \Big\{ (u^T\tilde{v}_i)^3 - (u^T\bar{v}_i)^3 \Big\} \\ &\leq \max_{\|u\|=1} \sum_{i=1}^r \frac{1}{\sqrt{q_i}} \Big(u^T(\widetilde{V} - \bar{V})e_i \Big) \Big((u^T\tilde{v}_i)^2 + (u^T\tilde{v}_i)(u^T\bar{v}_i) + (u^T\bar{v}_i)^2 \Big) \\ &\leq \frac{1}{\sqrt{q_{\min}}} \|\widetilde{V} - \bar{V}\|_2 (3 + 6\varepsilon_M + \varepsilon_M^2) \\ &\leq \frac{13}{\sqrt{q_{\min}}} \varepsilon_M \ , \end{aligned}$$

where the last line holds for $\varepsilon_M \leq 1/2$. Since $\|\tilde{H} - \bar{G}\|_2 \leq \varepsilon_H + (13/\sqrt{q_{\min}})\varepsilon_M$, we show that $\hat{V}^{\tilde{H}}$ and $\bar{V}$ are close using the perturbation analysis for robust power method from [2].

Theorem C.2 (Restatement of Theorem 5.1 by [2]). *Let $G = \sum_{i\in[r]} \lambda_i(v_i \otimes v_i \otimes v_i) + E$, where $\|E\|_2 \leq C_1 \frac{\lambda_{min}}{r}$. Then the tensor power-method after $N \geq C_2(\log r + \log\log\left(\frac{\lambda_{max}}{\|E\|_2}\right)$, generates vectors $\hat{v}_i, 1 \leq i \leq r$, and $\hat{\lambda}_i, 1 \leq i \leq r$, s.t.,*

$$\|v_i - \hat{v}_{P(i)}\|_2 \leq 8\|E\|_2/\lambda_{P(i)}, \quad |\lambda_i - \hat{\lambda}_{P(i)}| \leq 5\|E\|_2. \tag{39}$$

where P is some permutation on $[r]$.

Applying the above theorem to $\tilde{H}$ and $\bar{G}$, we get that $\|\hat{V}^{\tilde{H}} - \bar{V}\|_2 \leq 8\sqrt{r\,q_{\max}}(\varepsilon_H + (13/\sqrt{q_{\min}})\varepsilon_M)$. Notice that to apply the perturbation analysis, it is crucial that we use the fact that $\bar{G}$ has an orthogonal decomposition. Similarly, we can show that

$$\Big|\hat{\lambda}_i - \frac{1}{\sqrt{q_i}}\Big| \;\leq\; 5\varepsilon_H + \frac{65}{\sqrt{q_{\min}}}\,\varepsilon_M\;.$$

C.2 Proof of Lemma B.2

Let $E = \mathcal{P}_{\Omega_2}(S_2 - \mathbb{E}[S_2]) = S_2 - \mathbb{E}[S_2] - \mathrm{diag}(S_2 - \mathbb{E}[S_2])$, and we bound each term separately using concentration inequalities. Define the random matrix $E^{(1)} \equiv S_2 - \mathbb{E}[S_2] = \frac{2}{|\mathcal{S}|}\sum_{t\in\{1,\ldots,|\mathcal{S}|/2\}}\Big(x_t x_t^T - \mathbb{E}[x_t x_t^T]\Big)$. We apply the following matrix Bernstein bound for sum of independent sub-exponential random matrices with $X_t = x_t x_t^T - \mathbb{E}[x_t x_t^T]$.

Theorem 7 (Theorem 6.2 of [24])**.** *Consider a finite sequence $\{X_t\}$ of independent, random, self-adjoint matrices with dimension N. Assume that*

$$\mathbb{E}[X_t] = 0 \quad \textit{and} \quad \mathbb{E}[X_t^k] \preceq \frac{k!}{2}R^{k-2}A_t^2 \quad \textit{for } k = 2,3,4,\ldots$$

Compute the variance parameter $\sigma^2 \equiv \|\sum_t A_t^2\|_2$. Then for all $a \geq 0$,

$$\mathbb{P}\Big(\big\|\sum_t X_t\big\|_2 \geq a\Big) \;\leq\; N\,\exp\Big(\frac{-a^2/2}{\sigma^2 + Ra}\Big)\;.$$

The random matrix we defined X_t is zero-mean, and satisfies $\mathbb{E}[X_t^k] \preceq \mathbb{E}[(x_t x_t^T)^k] = \ell^{k-1}\mathbb{E}[x_t x_t^T]$, which follows form the fact that $x_t^T x_t = \ell$ almost surely. We can prove the inequality $\mathbb{E}[X_t^k] \preceq \mathbb{E}[(x_t x_t^T)^k]$ via induction. Let $\bar{X} \equiv \mathbb{E}[x_t x_t^T]$. When $k = 1$, $\mathbb{E}[X_t] = 0 \preceq \mathbb{E}[(x_t x_t^T)]$, since $\bar{X}$ is a convex combination of positive semidefinite matrices. When $k = 2$, $\mathbb{E}[X_t^2] = \mathbb{E}[(x_t x_t^T)^2] - \bar{X}^2 \preceq \mathbb{E}[(x_t x_t^T)^2]$, since $\bar{X}$ is positive semidefinite. Suppose $0 \preceq \mathbb{E}[(x_t x_t^T)^k - X_t^k]$ for some $k \geq 1$, then this implies $\mathbb{E}[X_t\big((x_t x_t^T)^k - X_t^k\big)X_t] \succeq 0$. It follows that

$$\mathbb{E}\big[\,(x_t x_t^T)^{k+2} + \bar{X}(x_t x_t^T)^k\bar{X} - (x_t x_t^T)^{k+1}\bar{X} - \bar{X}(x_t x_t^T)^{k+1} - X_t^{k+2}\,\big] \succeq 0\;,$$

which implies

$$\begin{aligned}\mathbb{E}\big[\,(x_t x_t^T)^{k+2} - X_t^{k+2}\,\big] &\succeq \mathbb{E}\big[\,(x_t x_t^T)^{k+1}\bar{X} + \bar{X}(x_t x_t^T)^{k+1} - \bar{X}(x_t x_t^T)^k\bar{X}\,\big]\\ &\succeq \ell^{k-1}\bar{X}(2\ell\mathbb{I} - \bar{X})\bar{X}\\ &\succeq 0\;.\end{aligned}$$

This follows from the fact that $\|\bar{X}\| \leq \mathbb{E}[\|x_t x_t^T\|] = \ell$. By induction, this proves the desired claim that $\mathbb{E}[X_t^k] \preceq \mathbb{E}[(x_t x_t^T)^k] = \ell^{k-1}\mathbb{E}[x_t x_t^T]$, for all k. Then, the condition in Theorem 7 is satisfied with $R = \ell$, $A_t^2 = \ell\,\mathbb{E}[x_t x_t^T]$, and $\sigma^2 = (|\mathcal{S}|/2)\,\ell\,\|\mathbb{E}[x_t x_t^T]\|_2$, which gives

$$\mathbb{P}\Big(\frac{|\mathcal{S}|}{2}\big\|\,E^{(1)}\,\big\|_2 \geq a\Big) \;\leq\; N\exp\Big(\frac{-a^2/2}{|\mathcal{S}|\ell\|\mathbb{E}[x_t x_t^T]\|_2/2 + \ell a}\Big)\;.$$

When $|\mathcal{S}| \geq (2\ell\log(N/\delta))/\|\mathbb{E}[x_t x_t^T]\|_2$ as per our assumption, the first term in the denominator dominates for $a = \sqrt{2|\mathcal{S}|\,\ell\,\|\mathbb{E}[x_t x_t^T]\|_2\,\log(N/\delta)}$. This implies that with probability at least $1-\delta$,

$$\|E^{(1)}\|_2 \;\leq\; \sqrt{\frac{4\,\ell\,\|\mathbb{E}[x_t x_t^T]\|_2\,\log(N/\delta)}{|\mathcal{S}|}}\;,$$

and since $\|\mathbb{E}[x_t x_t^T]\|_2 = \|\frac{\ell(\ell-1)}{N(N-1)}\mathcal{P}_{\Omega_2}(M_2) + (\ell/N)\mathbb{I}\|_2 \leq (\ell^2/N^2)\sigma_1(M_2) + (\ell/N)$, this gives the desired bound.

Now let $E^{(2)} \equiv \mathrm{diag}(S_2 - \mathbb{E}[S_2])$, where each diagonal term is distributed as binomial distribution Binom$(|\mathcal{S}|/2, \ell/N)$. When $|\mathcal{S}| \geq (2N/\ell)\log(|\mathcal{S}|/\delta)$ standard Bernstein inequality gives

$$\|E^{(2)}\|_2 \;\leq\; \max_{i\in[N]} E_{ii}^{(2)} \;\leq\; \sqrt{\frac{4\ell\,\log(N/\delta)}{N\,|\mathcal{S}|}}\;,$$

for all $i \in [N]$ with probability at least $1-\delta$. Together we have the desired upper bound on $\|E^{(1)}+E^{(2)}\|_2$.

In the case of $\|E_i\|$, similar concentration of measure shows that $|E_{ij}| \leq \sqrt{\frac{8\,\ell^2\,\log(N/\delta)}{N^2\,|\mathcal{S}|}}$ with probability at least $1-\delta$ for all $j \neq i$. This gives

$$\|\mathcal{P}_{\Omega_2}(E)_i\| \quad \leq \quad \sqrt{\frac{8\,\ell^2\,\log(N/\delta)}{N\,|\mathcal{S}|}}\,.$$

C.3 Proof of Lemma B.3

Let $\hat{H}_{abc} = \frac{n(n-1)(n-2)}{\ell(\ell-1)(\ell-2)}\frac{2}{|\mathcal{S}|}\sum_{t=1+|\mathcal{S}|/2}^{|\mathcal{S}|} Y^t_{abc}$ where $Y^t_{abc} = \sum_{(i,j,k)\in\Omega_3} x_{t,i}x_{t,j}x_{t,k}\hat{Q}_{ia}\hat{Q}_{jb}\hat{Q}_{kc}$, and $\hat{Q} = \tilde{U}_{M_2}\tilde{\Sigma}_{M_2}^{-1/2}$. Then,

$$\begin{aligned} Y^t_{abc} \quad = \quad & \langle x_t,\hat{Q}_a\rangle\langle x_t,\hat{Q}_b\rangle\langle x_t,\hat{Q}_c\rangle - \langle x_t,\hat{Q}_a\rangle\sum_{i\in[N]}(x^2_{t,i}\hat{Q}_{ib}\hat{Q}_{ic}) - \langle x_t,\hat{Q}_b\rangle\sum_i(x^2_{t,i}\hat{Q}_{ia}\hat{Q}_{ic}) \\ & -\langle x_t,\hat{Q}_c\rangle\sum_i(x^2_{t,i}\hat{Q}_{ia}\hat{Q}_{ib}) + 2\sum_{i\in[N]} x^3_{t,i}\hat{Q}_{ia}\hat{Q}_{ib}\hat{Q}_{ic}\,. \end{aligned}$$

We claim that $|Y^t_{abc}| \leq 6\ell^3\mu_1^3 r^{3/2}N^{-3/2}(1-\varepsilon_M)^{-3/2}\sigma_r(M_2)^{-3/2}$. Since x_t has only ℓ non-zero entries and by incoherence of $\mu(\hat{M}_2) = \mu_1$, we get that $|\langle x_t,\hat{Q}_a\rangle| \leq \ell\mu_1\sqrt{r/(N(1-\varepsilon_M)\sigma_r(M_2))}$. Similarly, $|\sum_{i\in[N]}(x^2_{t,i}\hat{Q}_{ib}\hat{Q}_{ic})| \leq \ell\mu_1^2(r/N)(1-\varepsilon_M)^{-1}\sigma_r(M_2)^{-1}$ and $|\sum_{i\in[N]}(x^3_{t,i}\hat{Q}_{ia}\hat{Q}_{ib}\hat{Q}_{ic})| \leq \ell\mu_1^3(r/N)^{3/2}(1-\varepsilon_M)^{-3/2}\sigma_r(M_2)^{-3/2}$.

Applying Hoeffding's inequality to $\hat{H}_{abc}$, we get that

$$\left|\hat{H}_{abc} - H_{abc}\right| \quad \leq \quad \frac{48\,r^{3/2}\,\mu_1^3\,N^{3/2}}{\sigma_r(M_2)^{3/2}}\sqrt{\frac{\log(2/\delta)}{|\mathcal{S}|}}\,,$$

for $\varepsilon_M \leq 1/2$ with probability at least $1-\delta$, where $H_{abc} = \mathcal{P}_{\Omega_3}(M_3)[\tilde{U}_{M_2}\tilde{\Sigma}_{M_2}^{-1/2},[\tilde{U}_{M_2}\tilde{\Sigma}_{M_2}^{-1/2},[\tilde{U}_{M_2}\tilde{\Sigma}_{M_2}^{-1/2}]$.

D Proof of Theorem 4 for the error bound of RANKCENTRALITY

The proof builds on key technical Lemma from [19]. Recall that the comparison graph $G=([n],E)$ has $N=|E|$ pairs/edges and the transition matrix for a random walk on this graph G for mixture component $a, 1\leq a\leq r$ is $\tilde{p}^{(a)} = [\tilde{p}^{(a)}_{i,j}] \in [0,1]^{n\times n}$ which depends on $\tilde{P}_a$, the estimation of P_a obtained by the algorithm in phase 1.

The Markov chain $\tilde{p}^{(a)}$ is designed in such a way that if $\tilde{P}_a$ were indeed exactly equal to P_a, based on the true model parameters $\mathbf{w}^{(a)}$, then the following holds (easy to check): (a) Markov chain $\tilde{p}^{(a)}$ is irreducible and aperiodic as long as G is connected, (b) the stationary distribution $\tilde{\pi}^{(a)}$ is equal (proportional) to $\mathbf{w}^{(a)}$ (without loss of generality, we assume that $\mathbf{w}^{(a)}$ are such that they sum up to 1). The above fact primarily holds because in this ideal scenario the correspond Markov chain is a reversible Markov chain with the desired stationary distribution. In reality, the Markov chain $\tilde{p}^{(a)}$ is an approximation of the ideal scenario and we need to quantity the error due to estimation error of $\tilde{P}_a$. This is precisely what we shall do next using the following Lemma of [19].

Lemma D.1 (Lemma 2 in [19])**.** *For a Markov chain $\tilde{p}$ and an aperiodic, irreducible and reversible Markov chain p with the stationary distribution π, let $\Delta = \tilde{p}-p$ and let $\tilde{\pi}^{(t)}$ be the distribution at time t according to the Markov chain $\tilde{p}$ when started with initial distribution $\tilde{\pi}^{(0)}$. Then,*

$$\frac{\|\tilde{\pi}^{(t)}-\pi\|}{\|\pi\|} \quad \leq \quad \rho^t\frac{\|\tilde{\pi}^{(0)}-\pi\|}{\|\pi\|}\sqrt{\frac{\pi_{\max}}{\pi_{\min}}} + \frac{1}{1-\rho}\|\Delta\|_2\sqrt{\frac{\pi_{\max}}{\pi_{\min}}}\,,$$

where $\pi_{\min} = \min_{i\in[n]}\pi_i$, $\pi_{\max} = \max_{i\in[n]}\pi_i$, $\rho = \lambda_{\max}(p) + \|\Delta\|_2\sqrt{\pi_{\max}/\pi_{\min}}$, and $\lambda_{\max}(p) = \max\{|\lambda_2(p)|,\ldots,|\lambda_n(p)|\}$ is the second largest eigenvalue of p in absolute value.

To prove the desired bound, we need control on the quantities $\|\Delta\|_2$ and ρ. We know that the spectral norm is bounded by $\|\Delta\|_2 \leq \varepsilon$ as per our assumption. The following lemma provides a lower bound on $1-\rho$.

Lemma D.2. *For $\varepsilon \leq (1/4)\xi b^{-5/2}(d_{\min}/d_{\max})$, we have*

$$1-\rho \quad \geq \quad \frac{1}{4}\frac{\xi}{b^2}\frac{d_{\min}}{d_{\max}} \; .$$

Note that we defined $b \equiv \pi_{\max}/\pi_{\min}$ and also it follows from the definition that $\|\pi\| \geq 1/\sqrt{n}$ and $\|\tilde{\pi}^{(0)} - \pi\| \leq 2$. Substituting the bounds on $\|\Delta\|_2 \leq \varepsilon$ and $1-\rho$, we get that there exists positive numerical constant C, C' such that for $t \geq C' \log\big(n/\varepsilon\big)/\log(\rho)$,

$$\frac{\|\tilde{\pi}^{(t)} - \pi\|}{\|\pi\|} \quad \leq \quad \frac{C\, b^{5/2}\, d_{\max}}{\xi\, d_{\min}}\, \varepsilon \; .$$

The necessary number of iterations can be further bounded by $t \geq C'(b^2 d_{\max}/(\xi d_{\min}))\big(\log(1/\varepsilon) + \log(n)\big)$ using Lemma D.2. This proves that the bound in Theorem 4, for $\varepsilon \leq (1/4)\xi b^{-5/2}(d_{\min}/d_{\max})$.

Now we are left to prove Lemma D.2. From the definition of ρ,

$$1-\rho \quad \geq \quad 1 - \lambda_{\max}(p) - \varepsilon\sqrt{b} \; .$$

We will show that

$$1 - \lambda_{\max}(p) \quad \geq \quad \frac{\xi d_{\min}}{2b^2 d_{\max}} \; , \tag{40}$$

where the desired bound follows for $\varepsilon \leq (1/4)\xi d_{\min}/(b^{5/2} d_{\max})$ as per our assumption. To prove (40), we use the comparison theorems (cf. see [19]), which bound the spectral gap of the Markov chain p of interest, by comparing it to a more tractable Markov chain. We use the simple random walk on the undirected graph $G([n], E)$ as a reference. Define the transition matrix of a simple random walk as

$$Q_{ij} \quad = \quad \begin{cases} \frac{1}{d_i} & \text{for } (i,j) \in E \; , \\ 0 & \text{otherwise} \; , \end{cases}$$

where d_i is the degree of node i. The stationary distribution of this Markov chain is $\mu_i = d_i / \sum_j d_j$, and Q is reversible since the detailed balance equation is satisfied, i.e. $\mu_i Q_{ij} = 1/(\sum_j d_j) = \mu_j Q_{ji}$ for all $(i,j) \in E$. The following key lemma provides a bound on the spectral gap of p with respect to the spectral gap of Q.

Lemma D.3 (Lemma 6 in [19])**.** *Let Q, μ and p, π be reversible Markov chains on a finite set $[n]$ representing random walks on a graph $G([n], E)$, i.e. $p_{ij} = 0$ and $Q_{ij} = 0$ if $(i,j) \notin E$. For $\alpha \equiv \min_{(i,j)\in E}\{\pi_i p_{ij}/(\mu_i Q_{ij})\}$ and $\beta \equiv \max_{i \in \mathcal{V}}\{\pi_i/\mu_i\}$,*

$$\frac{1-\lambda_{\max}(p)}{1-\lambda_{\max}(Q)} \quad \geq \quad \frac{\alpha}{\beta} \; .$$

We have defined $\xi \equiv 1 - \lambda_{\max}(Q)$, and α and β can be bounded as follows.

$$\begin{aligned} \alpha \quad &= \quad \min_{(i,j)\in E} \frac{\pi_i p_{ij}}{\mu_i Q_{ij}} \\ &\geq \quad \min_{(i,j)\in E} \frac{w_i w_j \sum_{k\in[n]} d_k}{d_{\max}(w_i + w_j)\sum_{k\in\mathcal{V}} w_k} \; , \text{ and} \\ \beta \quad &= \quad \max_{i\in[n]} \frac{\pi_i}{\mu_i} \\ &\leq \quad \max_{i\in[n]} \frac{w_i \sum_k d_k}{d_i \sum_k w_k} \; . \end{aligned}$$

Hence, $\alpha/\beta \geq d_{\min}/(2 d_{\max} b^2)$. This proves the desired bound in (40).

E Proof of Remark 2.1

Assuming $w_i^{(a)}$'s are drawn uniformly at random from the interval $[1, 2]$ and $G(\mathcal{V}, \mathcal{E})$ drawn from the Erdös-Rényi model with average degree $\bar{d} \geq \log n$, we want to bound the singular values and the incoherence of $M_2 = PQP^T$. Define a matrix $\tilde{M} = Q^{1/2}P^TPQ^{1/2} \in \mathbb{R}^{2\times 2}$. Since $\tilde{M}$ and M_2 have the same set of non-zero singular values, we analyze the spectrum of $\tilde{M}$.

For our example, $\tilde{M} = \frac{1}{2}P^TP$. Define P_1 and P_2 be the two columns of P such that $P = [P_1\, P_2]$, then using McDiarmid's inequality we get that, conditioned on the graph G with N edges and maximum degree $d_{\max}$,

$$\begin{aligned} \big|\, P_1^TP_2 \,\big| &\leq \varepsilon N\,, \\ \big|\, \|P_1\|^2 - (\ln(3486784401/68719476736)+3)N \,\big| &\leq \varepsilon N\,, \\ \big|\, \|P_2\|^2 - (\ln(3486784401/68719476736)+3)N \,\big| &\leq \varepsilon N\,. \end{aligned} \tag{41}$$

with high probability for any positive constant $\varepsilon > 0$. We provide a proof for $\|P_1\|^2$, and the others follow similarly. Conditioned on the graph G, $\|P_1\|^2 = \sum_{(i,j)\in\mathcal{E}}\big(\,(w_j^{(1)} - w_i^{(1)})/(w_i^{(1)} + w_j^{(1)})\,\big)^2 = f(w_1^{(1)}, \ldots, w_N^{(1)})$ is a function with bounded difference:

$$\sup_{w_1^{(1)}\ldots,w_n^{(1)},v_i^{(1)}} \big|\, f(w_1^{(1)},\ldots,w_{i-1}^{(1)},w_i^{(1)},w_{i+1}^{(1)},\ldots,w_n^{(1)}) - f(w_1^{(1)},\ldots,w_{i-1}^{(1)},v_i^{(1)},w_{i+1}^{(1)},\ldots,w_n^{(1)}) \,\big| \leq d_{\max}\,.$$

It follows that

$$\mathbb{P}\Big(\big|\, \|P_1\|^2 - \mathbb{E}[\|P_1\|^2] \,\big| \geq \varepsilon N \,\Big|\, G\Big) \;\leq\; 2\exp\Big\{-\frac{2\varepsilon^2N^2}{d_{\max}^2 n}\Big\}\,.$$

For Erdös-Rényi random graphs with $\bar{d} \geq \log n$, we know that $d_{\max} = \Theta(\bar{d})$ and $N = \Theta(\bar{d}N)$ with high probability. Also, it is not too difficult to compute $\mathbb{E}[\|P_1\|^2] = \ln(3486784401/68719476736) + 3 \simeq 0.0189$. It follows that with high probability, (41) holds.

Given (41), we can decompose the matrix as

$$\tilde{M} = \begin{bmatrix} (\ln(3486784401/68719476736)+3)N & 0 \\ 0 & (\ln(3486784401/68719476736)+3)N \end{bmatrix} + \Delta\,,$$

where $\|\Delta\|_2 \leq 2\varepsilon$. It follows that $\sigma_1(\tilde{M}) \leq 0.02N$ and $\sigma_2(\tilde{M}) \geq 0.017N$. Choosing $\varepsilon = 0.001N$, this proves the desired bound.

To bound the incoherence, consider the SVD of $M_2 = USU^T = PQP^T$. There exists a orthogonal matrix R such that $US^{1/2} = PQ^{1/2}R$. Then, the i-th row of U is $U_i = e_i^TPQ^{1/2}RS^{-1/2}$. We know, $Q = diag(1/2, 1/2)$, $S = \mathrm{diag}(\sigma_1(M_2), \sigma_2(M_2))$, and $RR^T = R^TR = \mathbb{I}$. It follows that

$$\begin{aligned} \mu(M_2) &= \max_i \sqrt{N/2}\|U_i\| \\ &\leq \max_i \sqrt{N/2}\sqrt{P_{i1}^2 + P_{i2}^2}(1/\sqrt{2})\sqrt{1/\sigma_2(M_2)} \\ &\leq 15\,. \end{aligned}$$